ON ETCHEMENDY ON TARSKI ON LOGICAL CONSEQUENCE

Wylie Breckenridge

In 1936 in "On the Concept of Logical Consequence" (OCLC) Alfred Tarski gave an analysis of logical consequence that remains at the centre of standard modern treatments. In 1990 in *The Concept of Logical Consequence* (CLC) John Etchemendy argued that Tarski's analysis is wrong. Many authors have criticised Etchemendy's arguments, but in his draft paper "Reflections on Consequence" (RC) he reaffirms and re-presents all of his main points, citing the "confusing, incomplete and otherwise misleading" way in which they were first presented as being the sole cause of much of this criticism. Etchemendy has convinced me. My aim in this essay is to bring together the main points that Etchemendy makes in these two works to show that Tarski's analysis is wrong, and to show that little can be done to correct it. I have broken the essay into five sections. In the first I state the aim of Tarski's analysis; in the second I describe it; in the third I give Etchemendy's reasons for why it fails; in the fourth I look at how it might be defended; and in the fifth I look at what Etchemendy thinks are the consequences for logic.

1. What was Tarski's Aim?

Tarski makes it quite clear in OCLC what he is trying to achieve. His aim is to define what it is for an argument to be logically valid, for its conclusion to be a logical consequence of its premises. But he points out that logical consequence is a pre-theoretical concept, one "whose introduction into the field of strict formal investigation was not a matter of arbitrary decision on the part of this or that investigator" and that in defining it effort must be made "to adhere to the common usage of the language of everyday life" (p. 409). Thus he sees the purpose of giving a definition not to be the introduction of a *new* concept, but to be the analysis of an *already existing* one. He talks about defining "the proper concept of consequence, which is close in essentials to the common concept" (p. 413). What are those essentials? What are the most significant features of our intuitive concept of logical consequence? He notes two. The first is that if a sentence, X, follows from (i.e. is a logical consequence of) a class of sentences, K, then "it can never happen that both the class K consists only of true sentences and the sentence X is false" (p. 414). The second is that whether or not X follows from the sentences in K depends only upon their logical structure, and not upon "knowledge of the objects to which the sentence X or the sentences of the class K refer" (p. 414-5). Tarski's aim in the paper is to give an analysis of logical consequence that agrees with at least these two features of our intuitive concept.

2. What was Tarski's Analysis?

Some terminology. Call an argument a *material consequence* if it preserves truth (i.e. does not have true premises and a false conclusion), a *consequence* if it necessarily preserves truth (i.e. *cannot* have true premises and a false conclusion), a *logical consequence* or *logically valid* or just *valid* if its conclusion is, intuitively, a logical consequence of its premises, and *logically invalid* or just *invalid* if it is not logically valid.

Consider this logically valid argument:

A1: 1. If Tarski was right then Etchemendy is wrong.
2. Tarski was right.
∴3. Etchemendy is wrong.

The two features of our concept of logical consequence that Tarski wants his analysis to capture are these: We think that the conclusion is a *consequence* of the premises because if the premises are true then the conclusion must also be true; and we think that it is a *logical* consequence because what makes it a consequence is its logical form and nothing else. It follows from A1 being a logical consequence, then, that a certain class of arguments must all be material consequences. Which arguments? Firstly, A1. Why must A1 be a material consequence? If it wasn't then it would have true premises and a false conclusion, so it would not be the case that it cannot have true premises and a false conclusion, so it would not be a consequence, so it would not be a logical consequence. Secondly, any argument of the same logical form as A1. Why must these all be material consequences? Suppose one of them wasn't. Then it would have true premises and a false conclusion, so it would not be the case that it cannot have true premises and a false conclusion, so it would not be a consequence, so it would be an argument of the same logical form as A1 but which is not a consequence, so A1 would not be a consequence *in virtue of its logical form*, so A1 would not be a *logical* consequence. What do we mean by 'the same logical form'? That's a question I will address soon. But it seems clear that whatever we mean, A1 should have the same logical form as itself. So we can say this: it follows from A1 being a logical consequence that every argument of the same logical form as A1 must be a material consequence. Or this: A1 is a logical consequence *only if* every argument of the same logical form as A1 is a material consequence. Or this: it is a necessary condition on A1 being a logical consequence that every argument of the same logical form as A1 is a material consequence. We can generalise: For *any* argument, Ω, let F(Ω) be the class of arguments of the same logical form as Ω (whatever that might mean). Then the following condition is a necessary condition on Ω being a logical consequence:

MC(Ω): Every argument in F(Ω) is a material consequence.

Tarski's claim is that there is a way of specifying F(Ω) - a way of defining what it means for two arguments have the same logical form - such that this becomes a *sufficient* condition as well, such that Ω is a logical consequence *if and only if* every argument in F(Ω) is a material consequence. This is Tarski's analysis. It remains to say what he thinks that specification of F(Ω) is.

As I said above, F(Ω) ought to include Ω itself - whatever we mean by 'the same logical form' we would want any argument to count as having the same logical form as itself. If that were all, if F(Ω) consisted of just Ω, then MC(Ω) would hardly be sufficient for Ω to be a logical consequence. For consider this argument:

A2: 1. Grass is red.
∴2. The Earth is flat.

A2 is a material consequence, because the premise is false. So if F(A2) consisted of just A2 then MC(A2) would be true - every argument in F(A2) would be a material consequence. So if MC(A2) was sufficient for A2 to be a logical consequence then A2 would be a logical consequence. But it's not.

F(Ω) ought to also include those arguments that can be obtained from Ω by replacing some (or none, or all) of its non-logical expressions with other expressions of the same grammatical type. For example, if we take the only logical expression in A1 to be "if ... then ...", then F(A1) ought to include the argument obtained by replacing the name "Tarski" by the name "Newton", and the argument obtained by replacing the predicate "is wrong" by the predicate "is hairy", and the argument obtained by replacing the proposition "Tarski was right" by the proposition "grass is green", and so on. That is, it ought to contain these arguments:

A3:
1. If Newton was right then Etchemendy is wrong.
2. Newton was right.
———
∴3. Etchemendy is wrong.

A4:
1. If Tarski was right then Etchemendy is hairy.
2. Tarski was right.
———
∴3. Etchemendy is hairy.

A5:
1. If grass is green then Etchemendy is wrong.
2. Grass is green.
———
∴3. Etchemendy is wrong.

Since we get these arguments by substituting into non-logical expressions while keeping the logical expressions fixed, then they all have, intuitively, the same logical form, so they all ought to be included in F(Ω). If F(Ω) included just these arguments and no more (that is, if F(Ω) was *defined to be* this class of arguments), would that make MC(Ω) a sufficient condition for Ω to be a logical consequence? No, because it would tie logical consequence too closely with the expressive resources of the language being used. To see why, consider this argument:

A6:
1. George Washington was president.
———
∴2. Abe Lincoln had a beard.

Intuitively, A6 contains no logical expressions and is not a logical consequence. Yet, if we were to define F(Ω) in the way suggested, and if we were to take MC(Ω) to be sufficient for Ω to be a logical consequence, and if we were to take all of the expressions of A6 to be non-logical, then it would be deemed a logical consequence in some languages. It would be deemed so, for example, in the language consisting of just two names, "George Washington" and "Abe Lincoln", and two predicates, "was president" and "had a beard". This because no matter how we interchange the names and predicates that appear in A6, the arguments that we obtain are all material consequences (none of them have a true premise and a false conclusion; in fact, none of them have a false conclusion). I should point out that the problem here is not that logical consequence would depend *in some way* upon the language being used. After all, whether or not an argument is a logical consequence ought to depend upon the meanings of the expressions that are marked out as 'logical': if "if ... then ..." had meant what "or" means, then A1 would no longer be (intuitively) a logical consequence. The problem here is that logical consequence would depend upon the *availability* of non-logical expressions in the language - we could turn an invalid argument into a valid one just by removing expressions from the language, or, going the other way, we could turn a valid argument into an invalid one just by adding expressions. That ought not to be the case.

Tarski suggests that we can remove this dependency on the language by 'looking through' its expressions directly to the world that they describe. His idea is that F(Ω) ought to include not those arguments that can be obtained by *substitutions* into the non-logical expressions of Ω, but by *interpretations* of them: interpretations, for example, of which individual a name refers to, which property (or class of individuals) a predicate refers to, and so on. To see how this removes the problematic dependency on language, consider argument A6 and the simple language described just after. F(A6) is now to include the argument obtained from A6 by interpreting "Abe Lincoln" as referring to Albert Einstein, and because Einstein did not have a beard this argument is not a material consequence, and hence A6 is (correctly) judged to not be a logical consequence. If we took *this* as the definition of F(Ω) would that make MC(Ω) sufficient for Ω to be a logical consequence? Tarski says *yes*. He claims that with F(Ω) defined in this way, MC(Ω) is a *necessary and sufficient* condition for Ω to be a logical consequence. That is his definition of F(Ω), and that is his analysis.

This is actually just a sketch of his analysis. To be thorough I would need to talk about, among other things, how the non-logical expressions of a language are sorted into grammatical categories and about what kinds of things the members of a category are thought of as referring to. But what I have said should be enough to bring out the following three important features. First, the analysis identifies the logical validity of an argument with the truth preservation of every argument in a certain class. That is, it is a *reductive* analysis - it reduces logical validity to truth preservation. Second, that class of arguments is obtained by interpreting the meaning of various expressions in various ways, *while keeping the world just as it is*. The world that the premises and conclusions of those arguments are taken to be making claims about is the *actual* world, not some other *possible* world. Third, it makes logical consequence *relative*: it is relative to a particular choice of logical expressions. The same argument may be deemed a logical consequence on one choice of logical expressions, but not a logical consequence on another. For example, A1 is deemed to be a logical consequence if we take "if ... then ..." to be its one and only logical expression, but it is deemed to *not* be if we take *none* of its expressions to be logical. This because the argument we get by interpreting "if ... then ..." as meaning "or", "Tarski was right" as meaning "Grass is green", and "Etchemendy is wrong" as meaning "Snow is black", is not a material consequence - it has true premises and a false conclusion. These features are important for the arguments to come.

What is Wrong with the Analysis?

Etchemendy argues that Tarski's analysis is both conceptually wrong (that is, it doesn't properly capture our concept of logical consequence) and extensionally wrong (that is, the class of arguments that it declares logically valid does not coincide with the class of arguments that actually *are* logically valid). I will look at each in turn.

First, the conceptual mistake. According to Tarski's analysis, to think that an argument form is logically valid is just to think that all of its instances preserve truth - there is *nothing more* to being logically valid than that. But there *is* something more, and it's easy to show. Consider the argument form *modus ponens* (of which A1 is an instance).

If the logical validity of modus ponens *consisted in* the truth preservation of all of its instances, then all that we could say of any instance whose premises are true is that *either* its conclusion is true *or* modus ponens is invalid. We could not deduce the truth of its conclusion from the truth of its premises, for to do so would require us to assume that modus ponens is valid, which would be to assume that all of its instances preserve truth, which would include assuming that this particular instance preserves truth, which would be to assume that its conclusion is true, which would be to assume the very thing that we want to deduce: we could not *deduce* the truth of the conclusion without *assuming* the truth of the conclusion. But clearly that is something that we can and do do. The logical validity of an argument form licenses us to argue from the truth of the premises to the truth of the conclusion. It gives us an *independent guarantee* - independent, that is, of the actual truth values of premises and conclusions - that all of the instances of the argument form preserve truth. It is this independent guarantee that enables us to establish the truth of the conclusion from the truth of the premises. It is this that gives logical validity its epistemic importance, and Tarski's analysis misses it. Etchemendy says: "the reductive analysis just omits the guarantee, attempting to replace it with that which the guarantee is a guarantee of" (RC p. 8). And: "It is like confusing the symptoms with the disease, effects with their cause" (p. 9). And: "All of the instances of *modus ponens* preserve truth because it is a logically valid argument form. This is true. What is false is that *modus ponens* is logically valid *simply because* its instances preserve truth" (p. 10).

Second, the extensional mistake. The *conceptual* error in Tarski's analysis does not guarantee that there will be an *extensional* error as well (after all, it would be conceptually wrong to define "has a heart" in terms of having a kidney, but it would give the right extension none the less), but neither does it guarantee that there won't be (it is just lucky that "has a heart" gives the right extension). The extensional adequacy is a further issue. Etchemendy argues that under certain conditions the analysis *will* get the extension wrong - under some it will overgenerate (that is, it will declare as logically valid some arguments that are not), under some it will undergenerate (that is, it will declare some logically valid arguments as not), and under some it will do both. Moreover, those conditions actually obtain.

The analysis will overgenerate, he argues, whenever the logical expressions of a language can make true non-logical claims about the world, and this will happen whenever the world is sufficiently homogeneous, or whenever the logical expressions of the language are sufficiently expressive, or both. To see the first, consider again the intuitively invalid argument A6 (George Washington was president, therefore Abe Lincoln had a beard). We saw that defining F(A6) in terms of substitution gets it wrong, because it depends on the expressiveness of the language. So Tarski defined it in terms of interpretation. But that makes it depend upon non-logical facts about the world instead. Why? A6 is judged by Tarski's analysis to be logically invalid because it *happens* to be the case that someone didn't have a beard (I chose Einstein). If it had just so happened that everyone had had a beard, then no matter how we interpret "Abe Lincoln", "Abe Lincoln had a beard" will be true, and A6 will be a material consequence in every case. No matter, we can get "Abe Lincoln had a beard" to come out false by interpreting "had a beard" as expressing some property that Abe Lincoln didn't have, or by simultaneously interpreting "Abe Lincoln" and "had a beard" so that the second refers to some property not possessed by the individual that the first refers to. But this still depends on a *non-logical* fact about the world - that there is at least one individual that lacks at least one property. Had the world been more homogeneous - had there been no individuals, for example - then A6 would have wrongly been judged a logical consequence.

To see the second - that sufficiently expressive logical expressions can lead to overgeneration - consider this argument:

A7:	1.	Grass is green.
	$\therefore$2.	There are at least twenty objects.

If we take our logical expressions to be the first-order quantifiers, the identity relation, and the truth-functional connectives, then the conclusion of this argument can be expressed using only logical expressions:

$$(\exists x_1)(\exists x_2)\ldots(\exists x_{20})((x_1 \neq x_2) \wedge (x_1 \neq x_3) \wedge \ldots \wedge (x_{19} \neq x_{20}))$$

So the logical expressions are expressive enough to make the non-logical claim that the world contains at least twenty objects. This is a problem for Tarski's analysis. Why? Because as it happens there *are* at least twenty objects in the world, so the conclusion is true. Moreover, it remains true no matter how we interpret its non-logical expressions, because it *has no* non-logical expressions. Thus, all the arguments in F(A7) have true conclusions and hence are material consequences, so according to Tarski A7 is logically valid. But intuitively it's not.

These two examples do not show that Tarski's analysis *does* overgenerate. The first shows that it *would* have done had the world been more homogeneous. But it's not. The second shows that it *would* do so were we to include the identity relation in the set of logical expressions. But nothing in the analysis requires us to do that, and if we don't then we have no problem. So let's leave it out. But these moves do nothing to assure us that it *never* overgenerates. Is the world sufficiently heterogeneous that *no* language can give rise to cases like the first? And if it turns out that our set of logical expressions is too expressive, can we *always* just leave one of them out?

No. In fact, leaving the identity relation out of the list of logical expressions would lead the analysis to *undergenerate* instead. It would, for example, declare the following intuitively valid argument invalid:

A8:	1.	Triangle(a)
	2.	a = b
	$\therefore$3.	Triangle(b)

Why would it be declared invalid? To be valid, this argument must preserve truth under all interpretations of its non-logical expressions - the predicate "Triangle", the constants "a" and "b", and the relation "=" (which we are taking to be non-logical here). But if we interpret "Triangle(a)" as meaning "a is a triangle", interpret "a = b" as meaning "a has (strictly) fewer sides than b", interpret "a" as referring to some particular triangle, and interpret "b" as referring to some particular quadrilateral, then we get an argument which does *not* preserve truth - it has true premises but a false conclusion.

In general, Tarski's analysis will undergenerate whenever there are intuitively valid arguments in a language whose validity depends, like A8, upon the meaning of expressions not contained in the language's list of logical expressions. This fact is already familiar: many intuitively valid arguments (such as "All men are mortal, Socrates is a man, therefore Socrates is mortal") are declared invalid in propositional logic. The problem that A7 and A8 together show is that there is a tension between overgeneration and undergeneration. If we include the identity predicate in the list of logical expressions, then the analysis overgenerates (it declares A7 to be logically valid). If we exclude it from the list then it undergenerates (it declares A8 to be invalid). And we can't get around the problem of overgeneration by excluding one of the other logical expressions instead, because that would lead to undergeneration elsewhere: excluding the existential quantifier would render invalid any argument whose validity depends upon the meaning of the existential quantifier (such as existential generalization), excluding the conjunction operator would render invalid any argument whose validity depends upon the meaning of the conjunction operator (such as simplification), and so on. One way we get overgeneration, the other way we get undergeneration. Either way the analysis gets the extension wrong. And this happens under conditions that *do* obtain. The analysis *is* extensionally incorrect.

3. <u>How Might the Analysis be Defended?</u>

In this section I will consider three of the most promising ways of defending Tarski's analysis and argue that each meets with little success, if any.

First, it can be shown, as Tarski points out (OCLC p. 417), that

(1) If the premises of an argument declared logically valid by his analysis are true then the conclusion *must* be true.

So it appears, contrary to my claim in the previous section, that Tarski's analysis *does* capture something of that independent guarantee that gives logical consequence its epistemic importance. But we must be careful here. (1) is ambiguous, because there are at least two ways of interpreting the scope of the word "must". If an argument is declared logically valid by Tarski's analysis (for *some* choice of logical expressions) and K is its set of premises and S is its conclusion, then say that S is a *Tarskian consequence* of K. Then on one interpretation the claim is this:

(1') The following three conditions are jointly incompatible:
► S is a Tarskian consequence of K;
► All the members of K are true;
► S is false.

On the other interpretation the claim is this:

(1") If S is a Tarskian consequence of K then the following two conditions are jointly incompatible:
► All the members of K are true;
► S is false.

It is the first interpretation, (1'), that can be shown. Indeed, its truth is almost trivial: Suppose S is a Tarskian consequence of K with all the members of K true and S false; to say that all the members of K are true and that S is false is to say that under *one* interpretation of the non-logical expressions (the intended interpretation) the argument has true premises and a false conclusion; but that contradicts the supposition that S is a Tarskian consequence of K - the supposition, that is, that under *no* interpretation of the non-logical expressions does the argument have true premises and a false conclusion. The second interpretation, (1"), is false. Indeed, its falsity is almost as trivial: Consider argument A6 (George Washington was president, therefore Abe Lincoln had a beard). There is a choice of logical expressions for which this is declared logically valid by Tarski's analysis - just choose *all* of the expressions to be logical, then because under the intended interpretation the argument preserves truth, and because there are no other interpretations allowed, the argument preserves truth under all interpretations, so is logically valid; thus "Abe Lincoln had a beard" is a Tarskian consequence of "George Washington was president"; so if (1") were true then "Abe Lincoln had a beard" and "George Washington was not president" would be incompatible; but they are not, so (1") is not true. Unfortunately for the defenders of Tarski's analysis, it is the truth of (1") that they need to even *begin* to argue that it captures some sense of the independent guarantee that is the important feature of our concept.

Second, there is a way of modifying Tarski's analysis so that it does *not* overgenerate in the case of first-order logic. The modification is to require that for Ω to be a logical consequence every argument in F(Ω) be truth preserving not only when its premises and conclusion are taken to be making claims about the world as a whole but also when they are taken to be making claims about every *subset* of the world when that subset is taken to *be* the world (call these subsets *domains*). Then we avoid the problem of overgeneration in the case of argument A7 without having to reduce our list of logical expressions and risk undergeneration. A7 comes out logically invalid on the modified account, even on the original choice of logical expressions, because there is an interpretation of its non-logical expressions and a domain for which it does not preserve truth: take that domain to be the first ten natural numbers, interpret "Grass" as referring to the number 2, and interpret "is green" as meaning "is even"; then the premise is true but the conclusion is false, so the argument does not preserve truth. It can be proved that the modified analysis does not overgenerate in the case of first-order logic. In fact, that is exactly what is done by any completeness theorem in any modern treatment of first-order logic (in all of which the modified analysis has now replaced Tarski's original as the standard). Some deductive system is defined whose rules of inference are intuitively logically valid, and which therefore generate only intuitively logically valid arguments. The completeness theorem then proves that all of the arguments declared logically valid by the modified analysis can be generated by the deductive system, and so are also intuitively logically valid. (Note: viewed this way, the completeness theorem uses the deductive system to assess Tarski's analysis, whereas it is usual to think of it as using Tarski's analysis to assess the deductive system.)

So the modified analysis is (provably) extensionally correct. But it feels *ad-hoc*, to have no principled motivation. Sure, it solves one problem with the original analysis, but it does so merely by treating one of its symptoms. It does nothing to stop the cause - it still provides no *conceptual* guarantee that the analysis will pick out only intuitively logically valid arguments, as it would if it genuinely captured our concept. It still relies, for example, on the world being a certain way to avoid overgeneration: if the world contained just 1000 individuals then this argument would be mistakenly declared logically valid:

A9: 1. Grass is green.

∴2. There are no more than 1000 individuals.

(The conclusion is true under all interpretations of its non-logical expressions (it has none) and for all domains.) And it does not stop the analysis overgenerating in the case of second-order logic. We can formulate sentences containing only the identity relation, truth-functions, and first- and second-order quantifiers that are true in all domains if and only if the Continuum Hypothesis is true, and other sentences that are true in all domains if and only if the Continuum Hypothesis is false. Let CH be a sentences of the first sort, and ¬CH be a sentence of the second sort. Then no matter whether the Continuum Hypothesis is true or false, one of these arguments will be declared logically valid:

A10: 1. Grass is green

∴2. CH

A11: 1. Grass is green

∴2. ¬CH

But neither is intuitively logically valid.[1] So even the modified analysis is extensionally incorrect for second-order logic.

Third, we can remove all of the analysis's dependency upon facts about the *actual* world in one fell swoop by requiring that for Ω to be a logical consequence every argument in F(Ω) be truth preserving not only when its premises and conclusion are taken to be making claims about the *actual* world but also when they are taken to be making claims about every *possible* world. This is a very attractive move. In fact, when Tarski talks about the "objects" that expressions are interpreted as referring to it is not immediately clear whether he means *actual* objects or *possible* objects. At one point he says that the substitutional approach to logical consequence (as opposed to his interpretational approach) "could be regarded as sufficient for the sentence X to follow from the class K only if the designation of all *possible* objects occurred in the language in question" (OCLC p. 416, my emphasis). By a possible object, here, does he mean an object in the actual world that it is possible to designate, or an object in some possible world? It's not clear. Either way, a moment's thought shows that defining logical consequence in this way would take away the very feature that makes the original analysis attractive. Suppose we define logical consequence in the way suggested. Which worlds do we need to count as *possible* in order for the analysis to pick out exactly the right class of arguments? To avoid *over*generation, we need as many worlds as we can get: the more worlds we have, the more chance we have of detecting invalid arguments - those whose truth preservation depends upon non-logical facts about the world. To avoid *under*generation, however, we can't allow too many worlds. We can't allow any world in which an intuitively logically true sentence is false. Why? Suppose that S is intuitively logically true (e.g. "it is either raining or it is not raining"), and that there is a

[1] To anyone who thinks that the continuum hypothesis is not a logical truth, that is. Anyone who thinks otherwise may not see a problem for Tarski's analysis here.

world, W, in which S is false; because S is intuitively logically true, the argument from no premises to S is intuitively logically valid; but it would not be declared so by the proposed analysis, because on the intended interpretation it does not preserve truth in world W. So the possible worlds that we need are all those worlds in which every logical truth is true - those we normally call *logically possible* worlds. Thus the definition of logical consequence would appeal to the notion of logical truth. But that would make it either useless or unnecessary. If we have no clear notion of logical truth, then the definition is useless. If we have a clear notion of logical truth then it is unnecessary - we could simply say that an argument is logically valid just in case the material implication from the conjunction of its premises to its conclusion is a logical truth (almost, anyway: this relies on the language having the material implication connective. Some tinkering is needed if it doesn't). The concepts of logical consequence and logical truth are so closely related that defining one in terms of the other is little help unless we also have a reductive analysis of the other.

5. What are the Consequences for Logic?

We have seen that the extensional adequacy of Tarski's analysis is sensitive to our choice of logical expressions. If we take them to be the truth-functional connectives, first-order quantifiers, and the identity relation, and if we use the modified analysis rather than the original, then we can prove (via soundness and completeness theorems) that we get exactly the right arguments coming out logically valid (thanks to the world being sufficiently heterogeneous). If we take any of these away then the analysis undergenerates: removing the identity relation renders A8 invalid, removing the existential quantifier renders existential generalisation invalid, removing the material conditional renders modus ponens invalid, and so on. If we add any more then the analysis can overgenerate. I have already mentioned that if we add the second-order quantifiers then one of A10 or A11 will wrongly come out logically valid. If we add "believes that" then the logical validity of

1. Peter believes that it is raining.

∴2. Someone else also believes that it is raining.

will depend upon whether or not any proposition is, as a matter of fact, believed by one and only one person. If there is no such proposition, as seems likely, then this argument will (counterintuitively) come out logically valid. Even if there is, then there is a value of n for which the following argument will definitely come out logically valid:

1. n people believe that it is raining.

∴2. n+1 people believe that it is raining.

Since there are only finitely many believers, we can take n to be any number greater than the number of believers. Then the premise cannot come out true on any interpretation of its non-logical expressions ("people", and "it is raining"). Etchemendy claims, and it is pretty clear, that it is not hard to find similarly problematic arguments if we add "it is necessary that", or "it is obligatory that", or "it has always been the case that", and so on, to our list of logical expressions.

The upshot of this is: if we regard Tarski's analysis as capturing our intuitive notion of logical consequence, then we cannot consider the second-order quantifiers, "believes that", "it is necessary that", and so on, as genuine logical expressions. Doing so would lead us into contradiction: there would be arguments which are *not* logically valid according to intuition but which *are* logically valid according to the analysis; but if Tarski is right then the analysis just *is* our intuition, so this would be a contradiction. Thus second-order 'logic', epistemic 'logic', modal 'logic', and so on (the 'logics' of second-order quantification, "believes that", "it is necessary that", and so on) would not really be logic at all. The only genuine logic would be first-order logic (with identity). But if Tarski is wrong - if his analysis does not capture our intuitive concept of logically consequence, and is known to get the extension wrong for some choices of logical expressions - then we are not bound to dismiss them. Etchemendy thinks that the recognition of Tarski's mistake ought to be welcomed by logicians as liberation.

REFERENCES

Etchemendy, John.

[CLC] *The Concept of Logical Consequence*, Harvard University Press, 1990. Reissued by CSLI Publications and Cambridge University Press, 1999.

[RC] “Reflections on Consequence”, Draft paper, 1999.

Tarski, Alfred.

[OCLC] “On the Concept of Logical Consequence”, 1936, in *Logic, Semantics, Metamathematics*, Oxford University Press, 1956.